Properties and Prediction

J. McKenzie Alexander

3 May, 2004

1 Introduction

It is a truism to say that the world admits description on different levels. Given this, however, an important task for philosophy of science (and for the philosophy of social science in particular) is to explain how descriptions offered at different levels are related. It is generally acknowledged that natural kind terms occuring in higher-level descriptions will not stand in any neat relation with natural kind terms occuring in lower-level descriptions. For example, economic theories contain natural kind terms such as "money" and "capital," neither of which are likely to be coextensive with the natural kind terms of physics or chemistry. If an exact relation between higher-level natural kinds and lower-level natural kinds exists, it will presumably be one of supervenience. Let us assume, then, that a natural kind term Tx occuring in a higher-level description is coextensive with some disjunction $S_1 x \vee S_2 x \vee \cdots \vee S_n x$ of natural kind terms occuring in lower-level descriptions.

It is also generally acknowledged (although perhaps less widely so than the previous point) that substituting coextensive expressions into laws may produce a formula which is not a law, even though it is a true general statement. For example, if '$T_1 x \to T_2 x$' is a higher-level law and the two natural kind terms '$T_1 x$' and '$T_2 x$' are coextensive with '$S_1\, x \vee S_2\, x \vee \cdots \vee S_n\, x$' and '$S_1'\, x \vee S_2'\, x \vee \cdots \vee S_m'\, x$', respectively, substituting the two disjuncts in for the natural kind terms produces:

$$S_1\, x \vee S_2\, x \vee \ldots \vee S_n\, x \to S_1'\, x \vee S_2'\, x \vee \ldots \vee S_m'\, x \tag{1}$$

which will hold in exactly the same conditions as the original law. Nevertheless, (1) does not seem likely to be a law: one can readily generate many such expressions having the form of (1), all of which will be true general expressions. However these general expressions, although true, do not carve nature at the joints in the way required of laws.[1]

Yet even if one accepts both of these claims, it does not seem that they generate any great difficulty for making predictions *across* levels of description. Suppose that one has a complete description of the world at some lower level, that one has full knowledge of the dynamical laws operating at the lower level, and that one has full knowledge of the relevant “bridge laws” connecting terms in the lower-level description with terms in the higher-level description. On the face of it, nothing precludes one from making predictions at the higher-level of description: given the complete description of the world at the lower level, apply the dynamical laws to obtain a prediction of the future state of the world at the lower-level, and then use the bridge laws to translate the description of the world from the lower-level to the higher level. Whether the resulting prediction at the higher-level will be exact or probabilistic depends upon the nature of the underlying dynamical laws and the bridge laws. However, the important point is that prediction across levels of description as outlined above seems straightforward and generates (I suspect) many of the motivating intuitions for reductionist programs. After all, even if formulae such as (1) are only true general statements rather than “laws,” for purposes of prediction true general statements suffice.

In what follows, I present a case in which prediction across two levels of description is *impossible*, in general. The reason for the impossibility has nothing to do with questions of meaning, coextensiveness of terms, or the nature of law. Rather, it derives from a curious interaction between our concept of prediction and the flexibility of the dynamics at the lower level of description. If correct, this result may have disturbing implications for our understanding of the relation between micro and macro levels of description that have not yet been considered.

2 A Simple Example

Consider the following scenario: a tribe of people, known as the Gavagai, live along a river. They are an peaceful, occasionally nomadic people who survive primarily by agriculture; as such, it is important that they have access to fresh water from the river and land on which to grow food. As a peaceful people, they do not engage in warfare with themselves or neighbouring tribes, and this, combined with the regular supply of food produced from their excellent agricultural skills, means that their population grows at a steady rate.

However, the Gavagai have little social structure. Each household is essentially an independent unit. As the population grows, whether each successive generation continues to live at the location they inherited from their parents depends on only one thing: the amount of available land around them. If there is enough land available, the children of a family will continue to live in the same location as their parents. If there is not enough land available, the children will relocate to some place along the river bank where available land exists. Gavagai families who need to relocate, though, do not relocate at random. Each family realizes that it is better to move upstream rather than downstream, as moving upstream insures that they have access to fresh water for irrigation. Gavagai families never move downstream.

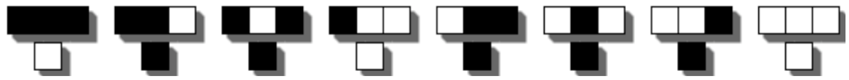

Figure 1: Migration patterns of the Gavagai.

We may represent the migration patterns of the Gavagai as follows: the river is represented as a horizontal line and a possible location of a Gavagai household as a "cell" along the line. If the location is occupied by a Gavagai household, we color the cell black; if the location is unoccupied, we color the cell white. If we adopt the convention that "upstream" is to the left and "downstream" is to the right, then the movements of Gavagai households occur as illustrated in figure 1. In that figure, the top line lists all of the possible inhabitation patterns for three adjacent blocks of land. The bottom line indicates the state of the center block of land in the next generation (a black square indicates that block of land will be inhabited by a Gavagai household, a white square indicates that block of land will be unihabited). One can verify that the migration patterns illustrated agree with our description: when three Gavagai households are together, the center household will leave in the next generation in search of more land. If the center household has available land on either side, it will remain in that spot in the next generation. If a household has available land upstream (to the left), the children of that household will "colonize" that area when they are able to; however, no children of any household ever move downstream from where they were born.

Given this rule and a description of the current locations of Gavagai households, we may calculate the future locations of Gavagai households at any time. (Of course, in the real world, the Gavagai will eventually run out of places to move upstream. Let us set aside this complication for the moment; we will return to it in a later section.) The dynamical law governing the movement of households is purely deterministic, and the state of the population is completely known at all times. Let us consider this the lower-level description—a kind of purely material "social physics." Figure 2 illustrates 400 generations of the Gavagai's migration patterns.

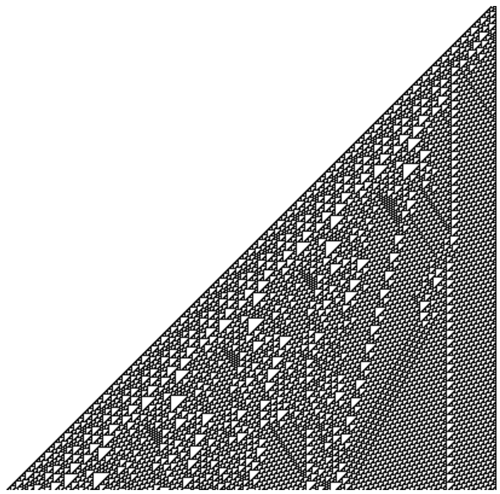

Figure 2: The migration patterns of the Gavagai — 400 generations starting from the initial state vector (1,0,1,1,1), where "1" (black) indicates occupation of territory and "0" (white) indicate empty space.

A sociologist studying the Gavagai may be interested in identifying certain states of societies as having certain properties. If we take a purely extensional view of properties, we may identify the properties studied by the Gavagai sociologists as simply sets of social states, where a social state is simply the distribution of Gavagai households along the river bank. This conception of a sociological property is, admittedly, a far cry from the complex ones occuring in sociological studies of real societies. My point is precisely that it is even with such a thin conception of a sociological property that impossible problems concerning prediction arise. Whether more complex conceptions of sociological properties (or higher-level properties in general) solve the problem or make matters worse I will not attempt to address here.

If a social property of the Gavagai is just a set of social states, how many social properties exist? According to a purely extensional conception of property, I cannot see a principled reason for excluding *any* set of social states from being considered a property. The set of properties of possible interest to a Gavagai sociologist, then, consists of all possible subsets of social states.

The above presents a very simple model in which descriptions of the world occur at two different levels. The lower-level description consists of a specification of the location of Gavagai households along the river bank; the higher-level description consists of attributions of various properties to the entire Gavagai society, where a social state S has property P simply if S is a member of P. The question of making predictions across levels of description can then be posed in the following way: given a complete description of the state of the Gavagai society at some time t (at the lower level), can the sociologist predict whether the Gavagai society will have property P at any future time, for any property P?

It is important to specify exactly what we mean by *predict*. I will assume that prediction, as used above, means that some effective procedure exists by which the sociologist can calculate whether the Gavagai society will have property P at any future time. This seems to me to capture the essence of prediction within a scientific theory. If no such procedure existed, one might be said to predict whether the Gavagai society will have property P at any future time, but the prediction would not be scientific. In this case, the predictor would be acting as a kind of divine oracle, capable of making pronouncements while unable to explain how he arrived at those pronouncements.

If we adopt this sense of prediction, it turns out to be impossible for a sociologist to be able to predict, for any property P, whether the Gavagai society will have P at some future time.[2]

3 The Argument

It turns out that that the migration patterns of Gavagai households, as illustrated in figure 1, are capable of universal computation. The transition rule illustrated in Figure 1, known as rule 110 in the cellular automata literature, was proven by Matthew Cook to be capable of universal computation.[3] We may use this to show that if the Gavagai sociologist were capable of predicting whether the Gavagai society would have property P at some future time, for any property P, then the Gavagai sociologist would be able to solve the halting problem (albeit in a rather circuitious fashion). Since the halting problem has no solution, the predictive capacity of the Gavagai sociologist must be limited.

Let M be a universal Turing machine consisting of a single infinite tape with cells indexed by the integers. We may assume that M begins with the read/write head initially positioned at 0 and that the inputs to M are located to the right (in cells indexed by positive integers), encoded in some format (the exact form does not matter). We may also assume that any output of M appears in cells to the left of the start position (in the cells indexed by negative integers) followed by a unique "end of computation" symbol, which M writes just before it halts. If M does not halt, then nothing will ever be written in the region of the tape having negative indexes; if M does halt, even if there is no actual output from the computation, at least the "end of computation" symbol will be written in the output region of the tape.

It is well-known that there exist many alternate models of computation equivalent to Turing machines. One can prove that there exist cyclic tag systems capable of emulating Turing machines.[4] Therefore, there is a cyclic tag system which emulates M. It can also be proven than any cyclic tag system can be emulated with rule 110 in a one-dimensional cellular automata, provided that the initial conditions are set appropriately. If we denote the entire nonempty input region for the universal Turing machine M by I, there is some corresponding set of initial conditions, denoted $\mathrm{Ic}(I)$, for a cellular automaton using rule 110 such that the forward evolution of the cellular automaton emulates M operating on I. That is, there is some initial distribution of Gavagai households along the river such that their migration patterns emulate a Turing machine operating on the given input.

Since the state of a cellular automaton is completely defined by the transition table and the state of the cells, it follows that if M produces any output from the input I (and halts), there is a set of patterns $P_{\mathrm{Ic}(I)} = \{P', P'', \ldots\}$ of cells representing this output. Because M only writes in the output region when M is about to end a computation, the appearance of any single pattern of $P_{\mathrm{Ic}(I)}$ indicates that M will halt on the input I.

Let $I_0, I_1, \ldots$ be an enumeration of the possible inputs to M. By choosing suitable encodings we can insure that $\bigcap_{i=0}^{\infty} P_{\mathrm{Ic}(I_j)} = \emptyset$, meaning that each set of patterns $P_{\mathrm{Ic}(I_i)}$ is unique. Let $P = \bigcup_{i=0}^{\infty} P_{\mathrm{Ic}(I_i)}$; that is, P is the set of all patterns of the cellular automaton representing the output of M when M halts on a given input.

Given this, we can now see why it is impossible for the Gavagai sociologist to be able to predict, given an initial distribution of Gavagai households along the river, whether any property will hold at some future time. The initial distribution of Gavagai households along the riverbank can be used to encode the initial conditions for a given input to a given Turing machine. If the Gavagai sociologist can predict whether any given property will hold of the Gavagai at some future time, the Gavagai sociologist could predict whether P would hold at some future time. That is, the Gavagai sociologist would be able to predict whether a Turing machine will halt on the input encoded by the initial distribution of Gavagai households along the river.

The actual problem is worse than the above discussion suggests. There are actually an infinite number of sets like P that preclude the sociologist's ability to make predictions, in general. Different Turing machines may expect their input to be encoded differently. By selecting different Turing machines $M', M'', \ldots$ and different encoding schemes $I_0', I_1', \ldots, I_0'', I_1'', \ldots$ for the inputs (and different cyclic tag systems emulating these Turing machines), an infinite number of sets $P', P'', \ldots$ exist, each of which may be used to solve the halting problem provided that the Gavagai sociologist was in possession of an effective procedure for prediction.

We can extend the constraints on the predictive capabilities of the Gavagai sociologist to concern sets of properties as well. Let $P_0, \ldots, P_n$ be a partition of P into a finite number of pieces. Clearly the sociologist cannot predict whether each of the properties $P_0, \ldots, P_n$ will hold at any future time because, if an effective procedure existed for *those* properties, we could construct an effective procedure for the original problematic property P. This can be said for each of the infinitely many problematic properties $P', P'', \ldots$ whose existence was argued for above.

Lastly, the predictive constraints can be extended to any (finite) set of properties such that by taking appropriate unions, intersections, or complements we can reconstruct P or any of the other problematic properties P', P'', $\ldots$. The result of this is that the class of properties for which the Gavagai sociologist cannot predict that they will hold at any future time is quite large, and that such properties need not have any obvious identifying features. Without serious reflection and a good understanding of the structure of sets of properties, we would be unable to say whether any randomly chosen property belonged to the set of problematic properties.

4 Objections and Implications

One objection to the problem posed above concerns the underlying conception of a property. Properties, I have assumed, are purely extensional and wildly prolific: any set of states is a property. Rejecting the extensional conception of a property will not resolve the problem for, even if properties are necessarily intentional, each property has an associated extension and this is all we need to develop the argument.

What if the set of actual properties is less prolific than I have assumed; that is, what if not every set of states corresponds to a property? Even if it is true that not every set of states corresponds to a property, this fails to remedy the situation. The crucial question concerns what properties *exist*, and, as long as we remain ignorant about the precise relation holding between higher-level properties and the lower-level states they supervene on—that is, which lower-level states constitute the higher-level properties—we just *do not know* whether P (or any of the other problematic properties) are among the existent properties (or are definable in terms of the existent properties).

A second objection notes that the problem encountered by the Gavagai sociologist could never arise in the *real* world because environmental constraints place an upper bound on the number of possible Gavagai and, even if no such upper bound existed, the Gavagai would eventually run out of available places to which to move. It is undeniably true that there is most likely an upper bound on the number of possible objects in the real world and the above argument requires (potentially) infinitely many. Yet should we be satisfied that the crucial fact settling whether we can, in general, make predictions across levels of description is an apparently arbitrary limit on the amount of matter in the universe? The problem encountered by the Gavagai sociologist raises important questions for our understanding of the connection between descriptions at the micro-level and predictions at the macro-level. Our understanding of why the problem arises or what it implies is not improved by noting that the alleged problem disappears if at most N things exist. More importantly, many scientific theories make use of infinitary arguments and allow for the existence of (possibly) infinitely many things. Since one does not need terribly complicated state transitions to construct the argument advanced above, I suspect that the problem of prediction across levels of description may occur in those theories just as easily as it does for the Gavagai sociologist.

A third objection may note that interesting predictions in the real world always have an implied upper bound on the duration of time involved. At the very least, we want to know whether a property will hold of some system before the end of the universe occurs. As with the second objection, this observation remains undoubtedly correct while neglecting the real point of the problem. Moreover, many predictions within scientific theories deliberately omit an implied upper bound on the time. When physicists ask whether the proton is stable, they are not asking whether the proton will decay in the next 5 years, or the next 500 years, or the next 500 billion years. They want to determine whether the proton will *ever* decay.

It seems, then, that the problem faced by the Gavagai sociologist is a potentially very real problem for our understanding of the relation holding between theories at different levels. Consider the relation between microeconomics and macroeconomics: given a low-level description of patterns of exchange between individuals in a market, can we use that description to encode computations, just like we used migration patterns of Gavagai households? If so, what implications does this have for our ability to predict macroeconomic properties from microeconomic descriptions? To say that it has none would be, at this point, unjustified. Similar questions can be raised for other social sciences.

I would like to close with two observations. First, if there are higher-level properties which we cannot predict from lower-level descriptions, for the reasons indicated in section 3, then we have a purely naturalistic and precise characterization of what one might call an *emergent* property. An emergent property, in this sense, has no mysterious metaphysical basis; it is, like all other properties, simply a set of states. Emergent properties *emerge* in the sense that we cannot always predict—indeed, it is *impossible* for us always to predict—whether they will ever hold at some future time. In this sense, properties P and P', P'', … are emergent properties. According to this definition, emergent properties may exist even for completely deterministic systems in which the state may be perfectly known at all times and for which, given a state description, we can predict what the state will be *at any time in the future.*

Second, the conditions described in the main argument offer one possible explanation for our inability to make accurate and reliable predictions in the social sciences. If the lower-level dynamics of social systems are capable of performing arbitrary computations, predicting the future state of society becomes equivalent to predicting the outcome of a computational process. For some computational processes, this is easy; for others, this is impossible. We will not be able to tell which outcome obtains without a better understanding of the dynamics of society.

Notes

[1] For the classic statement of this view, see Fodor's "Special Sciences, or the Disunity of Science as a Working Hypothesis," in *Synthese*, **28** (1974): 97–115.

[2] It is important to appreciate the order of quantification here. I am not arguing that, for any property P, it is impossible for the sociologist to be able to predict whether P will hold of the Gavagai society at some future time. There certainly are properties which the Gavagai sociologist will be able to predict. However, no general method (or methods) of prediction exist which enable the Gavagai sociologist to be able to predict whether an arbitrarily chosen property will hold.

[3] Matthew Cook never published his proof due to legal action by Stephen Wolfram, who claimed that publication would be in breach of agreements Cook entered into while employed by Wolfram. (See Giles, "What kind of science is this?" *Nature*, **417** (May 2002): 218.) Support for the claim that rule 110 is capable of universal computation can be found in Wolfram's *A New Kind of Science* (Champaign, IL: Wolfram Media, Inc., 2002), pp. 675–689 and pp. 1115–1116.

[4] I do not provide a definition of a cyclic tag system because the specific nature of such systems is essentially irrelevant for my argument. Wolfram's *A New Kind of Science* provides a reasonable introduction to cyclic tag systems.